Hardware Architectures for Ray-Tracing of Volume Data

Hanspeter Pfister

MERL - Mitsubishi Electric Research Laboratories

201 Broadway, Cambridge, MA 02139

Email: pfister@merl.com

SIGGRAPH 2001 EXPLORE INTERACTION AND DIGITAL IMAGES

Outline

VolumePro 500

- Traversal order and rendering pipeline.

RAYA

- Ray-Tracing Cache Design

SIGGRAPH 2001 EXPLORE INTERACTION AND DIGITAL IMAGES

VolumePro 500

First single-chip real-time volume rendering engine for consumer PCs.

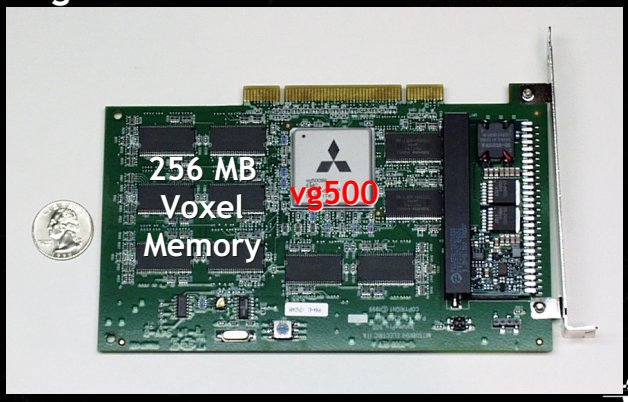

SIGGRAPH 2001 EXPLORE INTERACTION AND DIGITAL IMAGES

Special-Purpose Hardware

Research Projects

- VIRIM - University of Mannheim, Germany
- VOGUE - University of Tübingen, Germany
- VIZARD - University of Tübingen, Germany

VolumePro 500

- Cube-4 - SUNY Stony Brook, USA
- VolumePro 500 - Real Time Visualization

SIGGRAPH 2001 EXPLORE INTERACTION AND DIGITAL IMAGES

Volume Ray-Casting

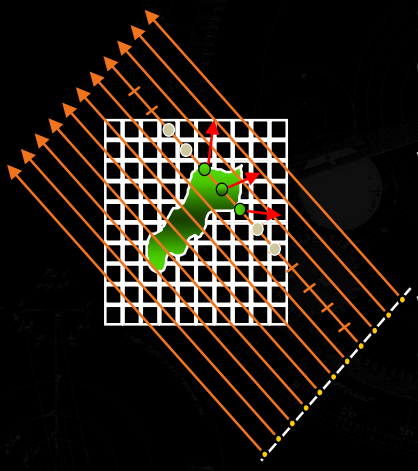

- Shoot rays for each pixel on the image plane.
- Best image quality of all volume rendering methods.
  - Easily accommodates super-sampling.
- Well understood.
- Simple implementation is not very efficient.

SIGGRAPH 2001 EXPLORE INTERACTION AND DIGITAL IMAGES

Image-Order Ray-Casting [Levoy 88]

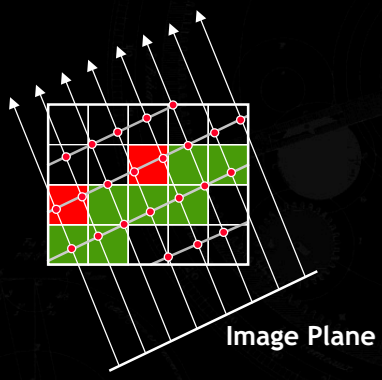

- Cast rays from image plane pixels.
- Irregular access to 3D tri-linear cells.
- No one-to-one mapping of samples to cells.
- Bad memory coherence.

SIGGRAPH 2001 EXPLORE INTERACTION AND DIGITAL IMAGES

Object-Order Ray-Casting

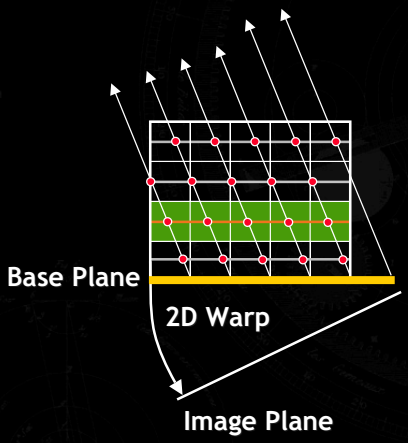

[Schroeder et al. 92]
[Yagel & Kaufman 92]
[Lacroute & Levoy 94]

Cast rays from base plane.

Traverse data on slices parallel to the base plane.

Final 2D warp.

Memory access coherence.

No perspective projections.

VolumePro System

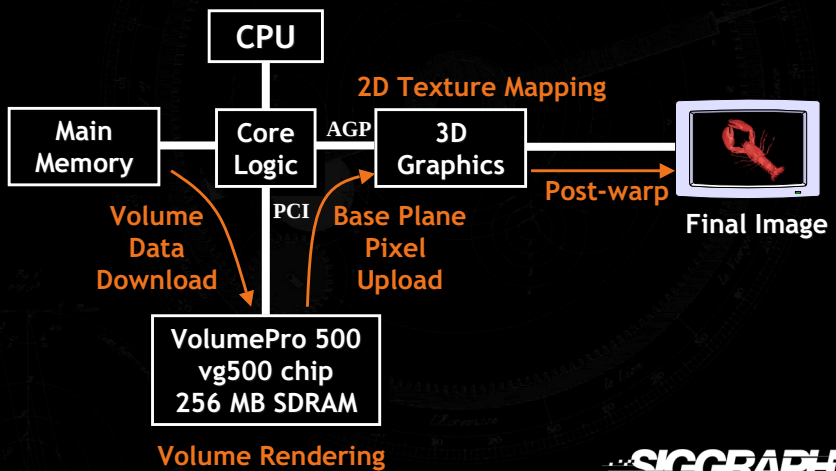

The vg500 Rendering Pipeline

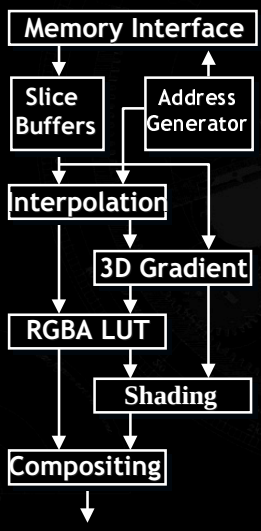

Full ray-casting pipeline, parallel projections only.

Fully pipelined at 125 MHz.

8- and 12-bit voxels.

Per-sample Phong shading.

Real-time classification.

500 million samples / sec.

- 30 Hz for 256^3 volumes.

Lessons Learned

Coherent memory access is a winner.

Brute force is beautiful.

Limited flexibility (brute force stinks).

Customers always want "one more feature".

- Integration with surface rendering.

Volume data doubles in size every 3 years.

- vg1000: 1 GB RAM, 512^3 @ 30 Hz.

Outline

VolumePro 500

- Traversal order and rendering pipeline.

RAYA

- Ray-Tracing Cache Design

RAYA Inspiration

"Rendering Complex Scenes with Memory-Coherent Ray Tracing"

Matt Pharr, Craig Kolb, Reid Gershbein, and Pat Hanrahan, Stanford University, SIGGRAPH '97

Ray Tracing of Volumes

Extending ray casting of volumes to ray tracing of volumes and geometry.

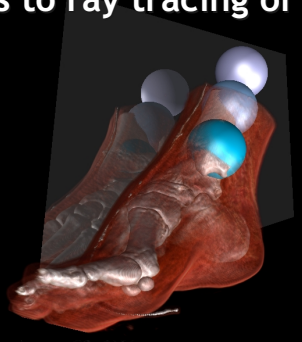

RAYA Assumptions

- Programmable rendering engine.
- Processor performance >> memory BW.
- Volume data >> on-board memory.
- Intelligent scene management.

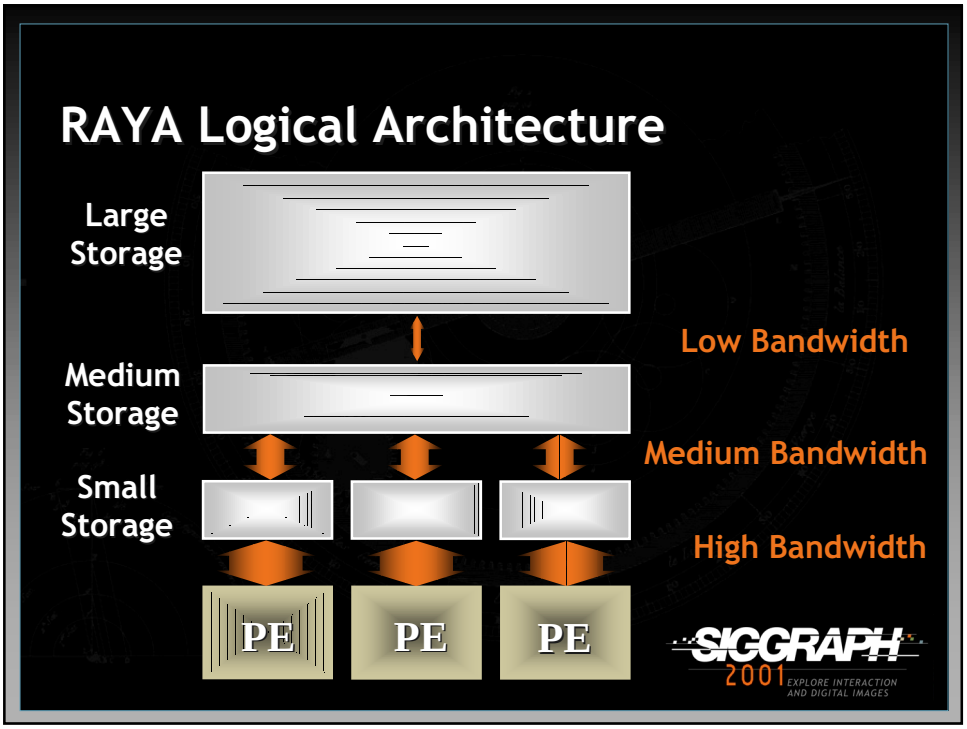
RAYA Logical Architecture
Large Storage
Medium Storage
Small Storage
PE
PE
PE
Low Bandwidth
Medium Bandwidth
High Bandwidth
SIGGRAPH 2001
EXPLORE INTERACTION AND DIGITAL IMAGES

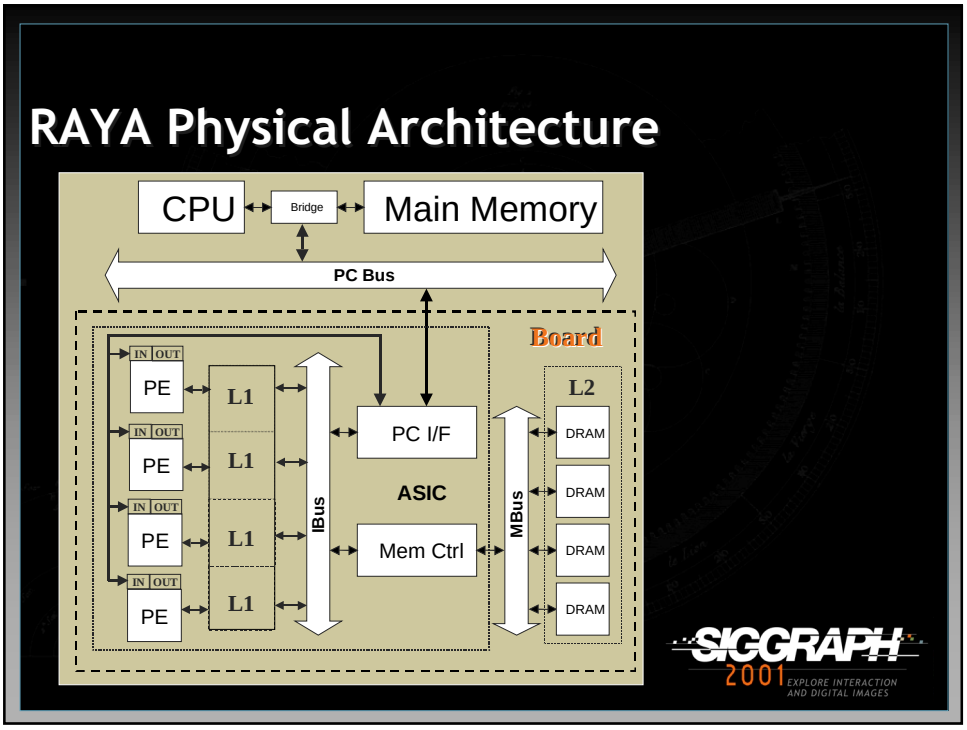
RAYA Physical Architecture
CPU
Bridge
Main Memory
PC Bus
Board
IN OUT
PE
L1
IBus
PC I/F
ASIC
Mem Ctrl
MBus
L2
DRAM
SIGGRAPH 2001
EXPLORE INTERACTION AND DIGITAL IMAGES

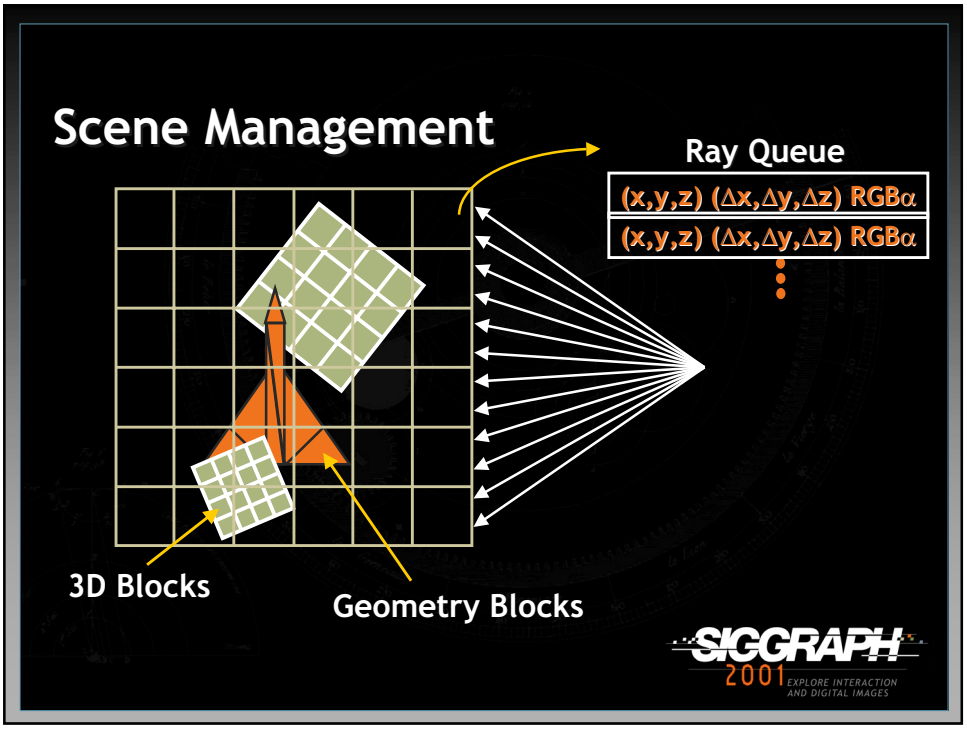
Scene Management
Ray Queue
(x,y,z) (Δx,Δy,Δz) RGBα
(x,y,z) (Δx,Δy,Δz) RGBα
3D Blocks
Geometry Blocks
SIGGRAPH
2001
EXPLORE INTERACTION AND DIGITAL IMAGES

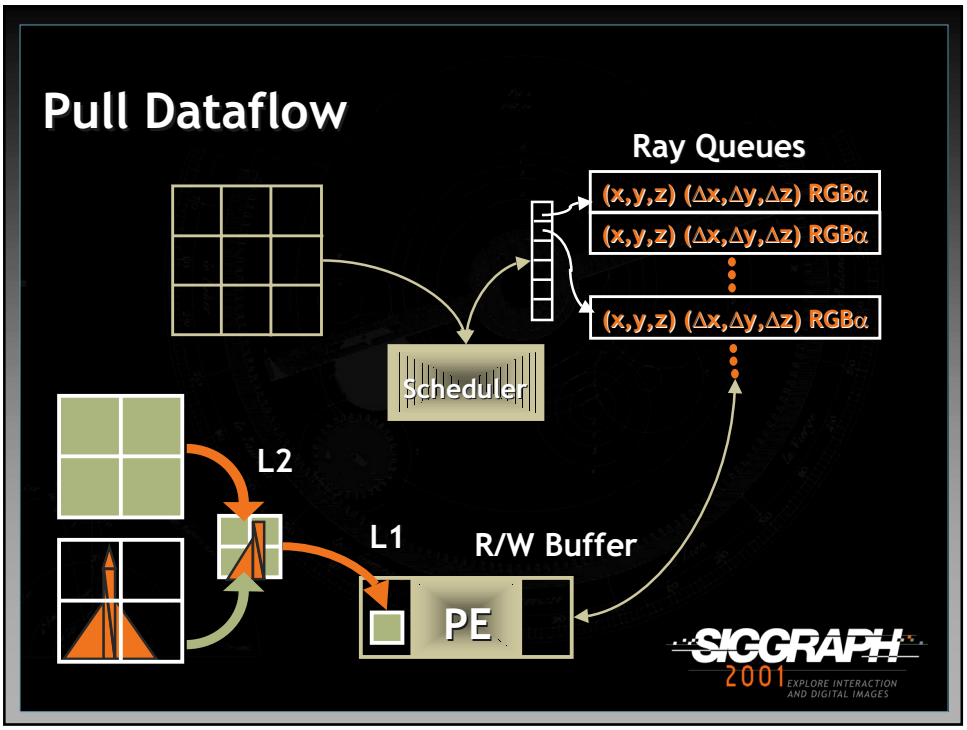
Pull Dataflow
Ray Queues
(x,y,z) (Δx,Δy,Δz) RGBα
(x,y,z) (Δx,Δy,Δz) RGBα
(x,y,z) (Δx,Δy,Δz) RGBα
Scheduler
L2
L1
R/W Buffer
PE
SIGGRAPH
2001
EXPLORE INTERACTION AND DIGITAL IMAGES

Advantages Pull

- L1 / L2 caches work independently.
- Simple PE algorithms.
- Graceful degradation if data does not fit in caches.
- Scheduling algorithm has less influence on performance.
- Simpler hardware architecture.

SIGGRAPH 2001 EXPLORE INTERACTION AND DIGITAL IMAGES

Example Scenes

- Medical Visualization.
- Large Volumes.
- Animation.
- Games.

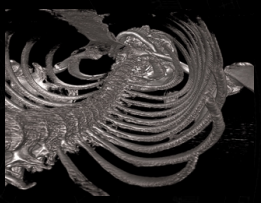

SIGGRAPH 2001 EXPLORE INTERACTION AND DIGITAL IMAGES

Results

L1 and L2 caching works very well.

Cache parameters:

- Line size 64 Bytes (L1), 512 Bytes (L2)
- Direct associative, 1 Mbit (L1), 32 Mbyte (L2)

Average Bytes Transferred / Frame

- 1.6 GB (L1), 45.1 MB (L2), 1.3 MB (MM)

Summary

Possible to build programmable RAYA with L1/L2 caches in a single chip.

Bottleneck is PE performance.

- How many GFLOPS?

Real-time volume ray-tracing is feasible for medium size volumes.

Acknowledgements

RTViz, Concord, MA
Erwin Janssen, TU Twente, Holland
Kevine Kreeger, SUNY Stony Brook, NY
Pat Hanrahan and his students, Stanford

SIGGRAPH 2001 EXPLORE INTERACTION AND DIGITAL IMAGES

Appendix

SIGGRAPH 2001 EXPLORE INTERACTION AND DIGITAL IMAGES

Working Sets

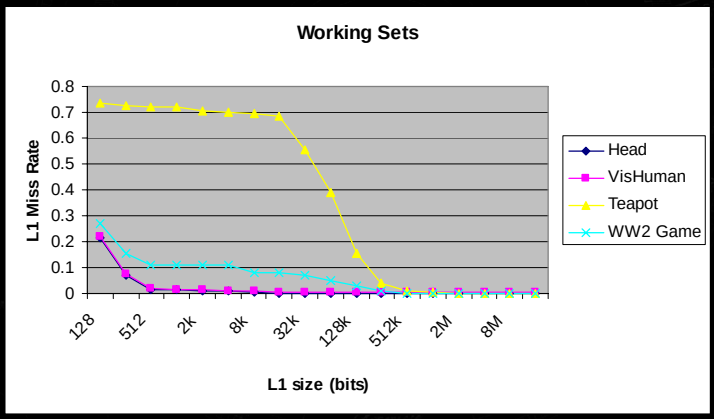

SIGGRAPH 2001 EXPLORE INTERACTION AND DIGITAL IMAGES

Volume, Geometry, Texture WS

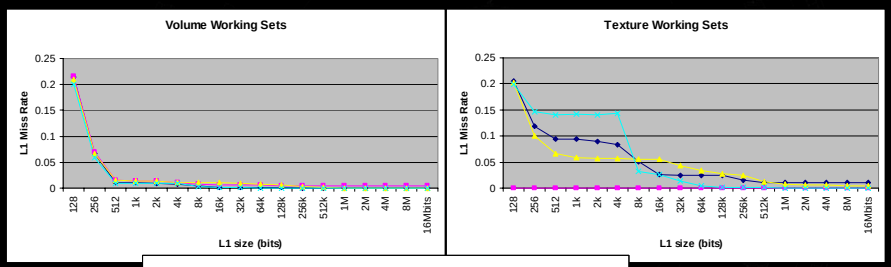

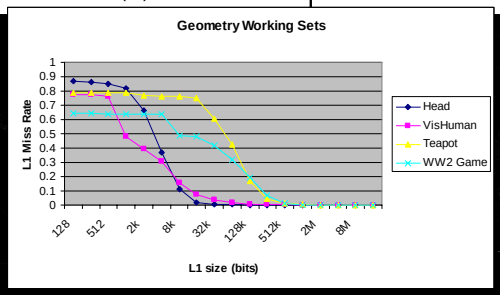

SIGGRAPH 2001 EXPLORE INTERACTION AND DIGITAL IMAGES

Associativity

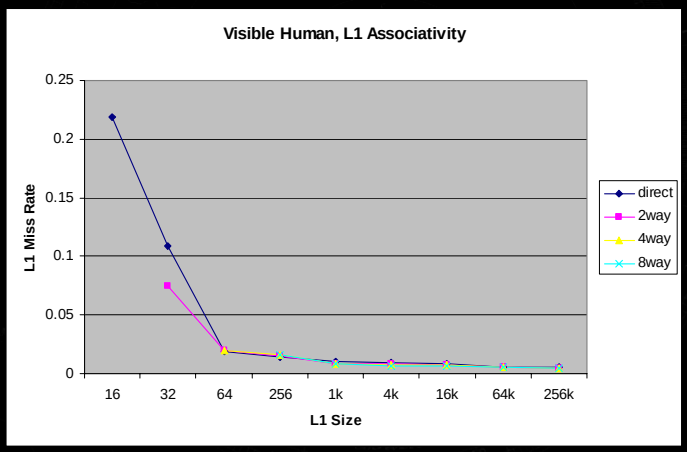

SIGGRAPH 2001 EXPLORE INTERACTION AND DIGITAL IMAGES

Average Bytes Read

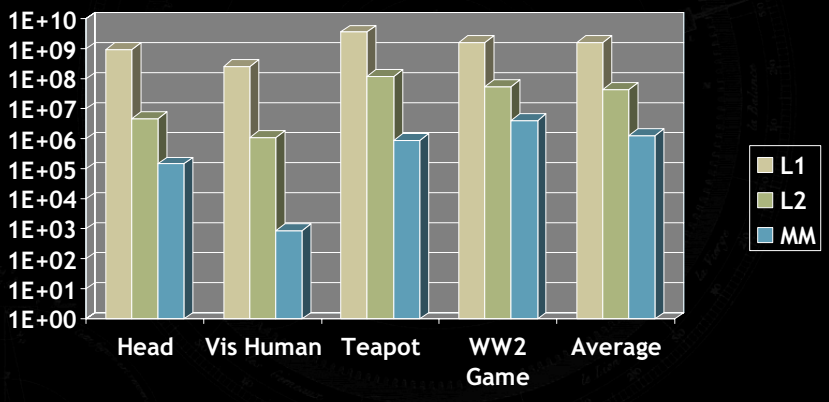

SIGGRAPH 2001 EXPLORE INTERACTION AND DIGITAL IMAGES